Available online at www.sciencedirect.com
SCIENCE DIRECT®
Nuclear Instruments and Methods in Physics Research A 544 (2005) 398–405
NUCLEAR INSTRUMENTS & METHODS IN PHYSICS RESEARCH Section A
www.elsevier.com/locate/nima

Simulation of drift compression for heavy-ion fusion

W.M. Sharp[a,*], J.J. Barnard[a], D.P. Grote[a], C.M. Celata[b], S.S. Yu[b]

[a]*Lawrence Livermore National Laboratory, L-440, Livermore, CA 94550, USA*
[b]*Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA*

Abstract

Lengthwise compression of space-charge-dominated beams is needed to obtain the high input power required for heavy-ion fusion. The "drift compression" scenario studied here first applies a head-to-tail velocity variation with the beam tail moving faster than the head. As the beam drifts, the longitudinal space-charge field slows compression, leaving the beam nearly monoenergetic as it enters the final-focus magnets. This paper presents initial work to model this compression scenario. Fluid and particle simulations are compared, and several strategies for setting up the compression schedule are discussed.

PACS: 52.58.Hm Heavy-ion inertial confinement; 52.65 Rr Particle-in-cell method

Keywords: Simulation; Drift compression

1. Introduction

Beams for heavy-ion fusion (HIF) must be compressed lengthwise by a factor of more than ten between an induction accelerator and the final-focus magnets. The compression scenario favored by the US HIF program is to impose a head-to-tail velocity increase or "tilt", so the beam tail approaches the head in a "drift-compression" section. The beam current and velocity must be accurately tailored or "shaped" before drift compression in order that the longitudinal space-charge field removes the velocity tilt just as the beam traverses the final-focus lattice. Transverse focusing in the drift-compression lattice must also be carefully designed to ensure that all parts of the beam remain approximately matched as the beam expands to the larger radius needed for final focusing.

An important problem posed by drift compression is how to prepare the beam velocity and current profiles before compression. Early work by Ho et al. [1] used analysis and one-dimensional numerical simulations to model a beam compression sequence that produced uniform current and velocity on

*Corresponding author. Fax: +1 925 424 6401.
E-mail address: sharp@hif.llnl.gov (W.M. Sharp).

0168-9002/$ - see front matter © 2005 Elsevier B.V. All rights reserved.
doi:10.1016/j.nima.2005.01.268

target. The effects of longitudinal space charge were ignored during profile shaping, on the grounds that shaping could be done in a sufficiently short time. A later paper by Sharp et al. [2] developed a more elaborate scheme where longitudinal fields in a beam-shaping section transform the input beam profile into the form found by allowing a compressed beam to expand backward through the drift-compression lattice. This shaping method was tested using a three-dimensional (3-D) fluid/envelope beam simulation that had a simple self-consistent model of the longitudinal space-charge field. Also, recent work by Qin et al. [3] uses an elegant three-dimensional envelope approach to develop several scenarios for self-similar compression.

This paper addresses two questions: whether the fluid models in earlier work adequately capture the physics of drift compression and what algorithm might be used to shape the beam current and velocity profiles of HIF beams before longitudinal compression. Other critical questions, such as how much the total emittance grows, whether a beam halo develops, and how these processes scale with beam and lattice parameters, will be addressed elsewhere. Section 2 describes three models of beam compression: a cold-fluid analytic description, a fluid-like simulation, and a particle simulation. In Section 3.1, we compare results of these models for an idealized compression scenario, and Section 3.2 uses the simulations to demonstrate several methods to prepare a beam for drift compression. We conclude with a brief summary of our findings.

2. Compression models

Three models are used here to describe beam compression. We compare the results of a simple cold-fluid analytic model with numerical simulations from a fluid/envelope code CIRCE [4] and from a 3-D particle-in-cell code WARP3d [5]. Two important measures of beam compression are the ratio of the initial length or duration to the value at stagnation, called the compression ratio, and the distance the beam midpoint propagates before stagnation, termed either the compression length or stagnation distance. In addition, the numerical models provide information about beam dynamics, such as the radial envelope and, for the WARP cases, transverse and longitudinal emittance.

With several simplifying assumptions, a simple cold-fluid model provides usable estimates of the compression ratio and length. The critical simplification is to model the longitudinal space-charge field E_z of the beam by what is called a "g-factor" model [6]. The transverse variation of E_z is ignored, and the dependence on the longitudinal coordinate z is given in SI units by

$$E_z = -(g/4\pi\varepsilon_0)\,\mathrm{d}\lambda_\mathrm{b}/\mathrm{d}z, \tag{1}$$

where λ_b is the beam line-charge density, ε_0 is the free-space permittivity, and g for a uniform-density beam with an average radius a centered in a beam pipe of radius R is

$$g = \ln(R^2/a^2). \tag{2}$$

In general, a will increase as the beam compresses, leading to a decrease in g. However, we assume that transverse focusing force increases appropriately with z, so that a, and hence g, remains fixed. For uniform density, the further assumption of a parabolic current variation in z leads to a space-charge field that varies linearly along the beam length. In the absence of longitudinal pressure, this linear field allows self-similar lengthwise compression of the profile. From the longitudinal envelope equation that results from these assumptions, we find that the compression ratio C is given by

$$C = 1 + (8\kappa_0 g)^{-1}(\Delta v/v)^2, \tag{3}$$

where κ_0 is the initial value of the generalized beam perveance, written for a non-relativistic beam with ion charge state q and mass M as

$$\kappa_0 = (4\pi\varepsilon_0)^{-1} 2qe\lambda_\mathrm{b}/(Mv_z^3), \tag{4}$$

and $\Delta v/v$, referred to here as the "velocity tilt," is shorthand for the initial difference between the longitudinal fluid velocity of the beam tail and that of the beam head, divided by the initial average longitudinal velocity v_z. The corresponding stagnation distance D is found to be

$$D/L_{\mathrm{b}0} = (\Delta v/v)^{-1}\{(1 - C^{-1}) + [(C-1)^{1/2}/(2C^{3/2})] \times \ln[2C + 2C^{1/2}(C-1)^{1/2} - 1]\}, \tag{5}$$

where we have scaled D to the initial beam length L_{b0}. By itself, the first term in Eq. (5) gives the propagation distance to reach minimum beam length for ballistic compression, while the second term corrects for space charge. For heavy-ion beams of interest for inertial fusion, typical perveance values range from 10^{-5} to 10^{-2}, and practical considerations are likely to limit the velocity tilt to less than 0.25 for driver-scale beams.

A second estimate of beam compression is obtained from numerical simulations using the code CIRCE [4], which divides a beam lengthwise into thin "slices," each containing a constant amount of charge. The transverse dynamics of each slice is calculated from envelope equations for the radii and the centroid position in the two transverse directions, and the longitudinal motion of each slice boundary is modeled by a Lagrangian fluid equation. The longitudinal space-charge field is approximated here by one of two methods: a generalized g-factor expression that accounts for variations of the beam density and the slice radii with z, and a Bessel series representation [6] that calculates the average E_z across slice boundaries from an exact analytic expression. The CIRCE dynamics model ignores longitudinal pressure, and it assumes that the transverse emittance of each slice is constant. Nonetheless, when benchmarked against the particle simulations, CIRCE is found to give reliable results, provided that the normalized emittance does not increase significantly and there is little longitudinal mixing. Because of its speed, CIRCE is a useful scoping code for lattice design.

We compare these two fluid-like models with results from the 3-D electrostatic particle-in-cell (PIC) code WARP3d [5]. Although this code can model lattice elements from first principles, we specify hard-edge focusing quadrupoles and fringe-free acceleration fields to keep the lattice representation as close to the CIRCE model as possible, and we initialize WARP simulations with matched, uniform-density, longitudinally cold beams, like those assumed in CIRCE. For this initial work, we choose that the beam, the lattice, and the fields are error free.

3. Results

For all the simulations presented here, we use parameters appropriate for the Integrated Beam Experiment (IBX), a scaled experiment being considered at LBNL that would test all the major accelerator components for HIF. Although the peak current for this accelerator would be less than 1 A, the tune depression would match that in a HIF driver, and with proper design, nearly all aspects of drift compression would be realistically modeled. The only physics that would not be correctly represented in IBX are image forces, stray-electron effects, and collisions between beam ions and with residual gas.

3.1. Comparison of compression models

To allow a fair comparison of the three models, we specify beams with an initially parabolic current profile and with uniform density and emittance, so that they match the assumptions of the analytic model. The ion species used here is singly charged potassium (39 amu) with an energy of 17 MeV, so the 6.9-A peak current gives a maximum perveance of about 8.8×10^{-4}. We limit the scale of the simulation by choosing a 1.75-m initial beam length, corresponding to a 250-ns duration. The transverse emittance, which is relatively unimportant in these simulations, is 4.3 mm mrad. A periodic FODO lattice is chosen to give an undepressed phase advance of 70° per 0.6-m lattice period and an initial average radius of 0.008 m.

The analytic compression model summarized by Eqs. (1) and (2) predicts a quadratic increase in the compression ratio with increasing initial velocity tilt $\Delta v/v$. Also, the compression length is predicted to increase from zero for $\Delta v/v = 0$ and drop off after a single maximum roughly like $(\Delta v/v)^{-1}$. The curves labeled "fixed g-factor" in Fig. 1 show these analytic relations for the chosen 8.8×10^{-4} perveance, and we find qualitatively similar curves for other perveance values.

Although the CIRCE simulations corroborate these qualitative predictions, we see significant quantitative differences in Fig. 1 between the various space-charge models. When we artificially

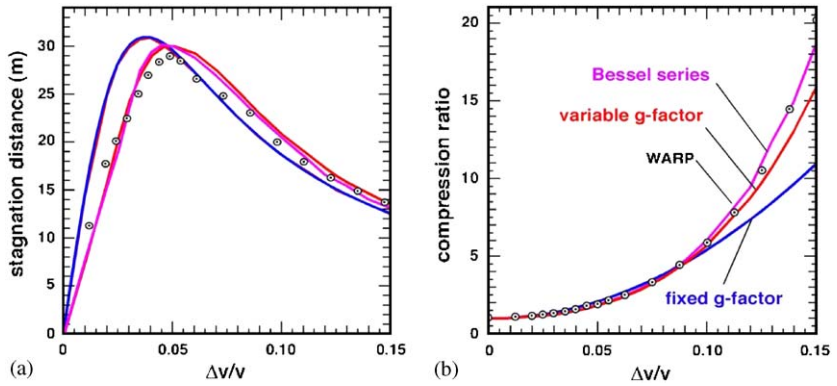

Fig. 1. Plots of (a) stagnation distance and (b) compression ratio for uniform-density parabolic IBX-like beams with a range of initial velocity tilts. The overlaid "fixed g-factor" curves are calculated from the analytic expressions and by CIRCE. The "variable g-factor" and "Bessel series" are calculated using CIRCE with more realistic models of the beam space-charge field. The points labeled "WARP" were calculated using WARP3d.

fix the g-factor in CIRCE, the code results accurately reproduce the analytic predictions for stagnation distance and compression ratio, validating the analysis. However, the more accurate space-charge models in CIRCE give significantly higher compression ratios when $\Delta v/v > 0.1$. Since the beam radius a is expected to increase as the beam compresses in a periodic focusing lattice, the g-factor in Eq. (2) and hence the space-charge field should decrease. Since it is this field that slows compression, the higher compression ratios seen for larger tilts are physically plausible, as are the longer stagnation distances seen in Fig. 1a for $\Delta v/v > 0.05$. For smaller initial tilts, the change in beam length and hence a becomes insignificant, and the stagnation distance is determined by how quickly the space-charge field at each end reverses the initial tilt. For a parabolic current profile, the fixed g-factor model is known to underestimate the longitudinal space-charge field near the beam ends [6], so for sufficiently small $\Delta v/v$, the model erroneously predicts longer stagnation lengths than CIRCE.

The corresponding WARP3d cases, plotted as points in Fig. 1, largely corroborate the CIRCE model, with agreement being best for the Bessel-series E_z representation. Generally, good agreement is seen in these cases because the transverse emittance in the WARP simulations grows less than 10%, and the longitudinal thermal energy remains small compared with the electrostatic potential energy, so two of the assumptions in the fluid/envelope model are approximately valid. The model is expected to fail for larger values of perveance and initial tilt, due to longitudinal and transverse heating, but numerical limitations in CIRCE prevent the exploration of these limits.

In addition to some randomness seen in the WARP3d stagnation distances, due to poor particle statistics near the beam ends, the results near the curve maximum lie systematically below the corresponding CIRCE values. The cause of this discrepancy is still under study, but it may result from two features of the WARP model missing from the fluid models. First, both the analysis and the CIRCE model average the longitudinal space-charge field radially, whereas the Bessel-series representation in Ref. [6] shows that the field is peaked on axis. Since it is this on-axis value that affects the velocity of the beam ends, a larger on-axis value should reduce compression and lead to earlier stagnation. A second difference is the absence of longitudinal pressure in the fluid models considered here. WARP plots of the longitudinal phase space show some broadening of the distribution in v_z as a beam compresses, and this higher longitudinal temperature resists compression. More careful analysis of WARP results is

needed to determine whether either of these effects is large enough to account for the observed reduction in compression length.

The WARP3d compression-ratio values plotted in Fig. 1(b) fall generally along the curve obtained from the Bessel-series CIRCE model. The one anomalous WARP result, seen for the highest $\Delta v/v$ value, is likely caused by particle loss at the beam-pipe wall during the last meter before stagnation, due to radial expansion as the beam is compressed. This loss, which approached 20% of the beam charge, may explain the larger compression ratio calculated for $\Delta v/v = 0.15$.

3.2. Comparison of compression strategies

A somewhat more realistic beam and lattice are used to compare methods for shaping the beam velocity and current profiles before drift compression. The beam is again 17-MeV potassium, but the 4-m beam length (1200-ns duration), 0.46-A peak current, and 10.3-mm mrad emittance match beam parameters from existing equipment that could be used as an IBX front end. The perveance for this beam is 8.4×10^{-4}, and the current profile has a 2-m flat top and drops to zero parabolically at the ends. A more complicated lattice is used for these cases. To prevent the outer edge of the beam from scraping on the beam-pipe wall, the field strength of focusing quadrupole magnets is gradually increased by about 35% over the second half of the 24-m lattice, and the beam-pipe radius increases stepwise from 0.03 m to 0.06 m over 12 lattice periods (7.2 m) beginning at 6.25 m. Also, 0.05-m induction gaps are placed midway between quadrupoles in the first 17 half-lattice periods to impose the beam-shaping electric fields.

The simplest method used here to impose a velocity tilt uses a previously published algorithm to control the beam energy and duration as functions of either time or the longitudinal position [7]. By assuming instantaneous acceleration in induction gaps and ignoring space-charge forces between gaps, we can trivially construct the trajectories of beam slices, and then calculate the corresponding acceleration fields as functions of time from the changes in slice trajectories at the gaps. Longitudinal control fields, called "ears," are added to the acceleration fields to counteract the effects of space charge while the tilt is being applied. We calculate these ear fields by running CIRCE with the longitudinal space-charge field artificially turned off. As the beam passes through each gap, the space-charge field is calculated as a function of time and used to estimate the required ears. For the case here, we choose that the beam duration decreases parabolically with time and the average energy increases sufficiently that all pulses are non-negative. The resulting shaping fields, shown in Fig. 2(a) as functions of time relative to the beam midpoint, are approximately triangular, and the corresponding ear fields, which retard ions at the beam head and slow those at the tail, reach a magnitude of about 4 kV at the beam ends. While this compression schedule achieves more than a ten-fold increase in the maximum current, as seen in Fig. 2(b), the compression is not self-similar. Since ear fields are applied only in the shaping section, no correction is made for space charge as the beam compresses. Consequently, at stagnation, nearly half the beam charge lies outside the pulse length at half-maximum, giving this case a compression ratio of about 5:1. As in the idealized runs of Section 3.1, the corresponding CIRCE run shows roughly the same compression, although the model overestimates the peak current by about 10%.

To control the beam-end expansion seen in Fig. 1, we use the algorithm in Ref. [7] to set up a shaping schedule that first imposes a velocity tilt in eight gaps placed a full lattice period apart, then corrects and finally removes the tilt in 13 more induction gaps spread over the remaining 19 m. Ears are added in all gaps, augmenting the shaping field as tilt is applied and reducing it as the tilt is removed. Simulations indicate that this approach largely eliminates blow-off of the beam ends, so that less than 12% of the beam charge in the WARP and CIRCE profiles is outside the full-length at half-maximum. However, the shaping method is unsatisfactory for two important reasons. Due to the wider spacing of gaps, the maximum field strength is nearly double that of Fig. 2(a), and the very short rise and fall times of the beam near stagnation require voltages for the ear pulses approaching 400 kV and roughly a

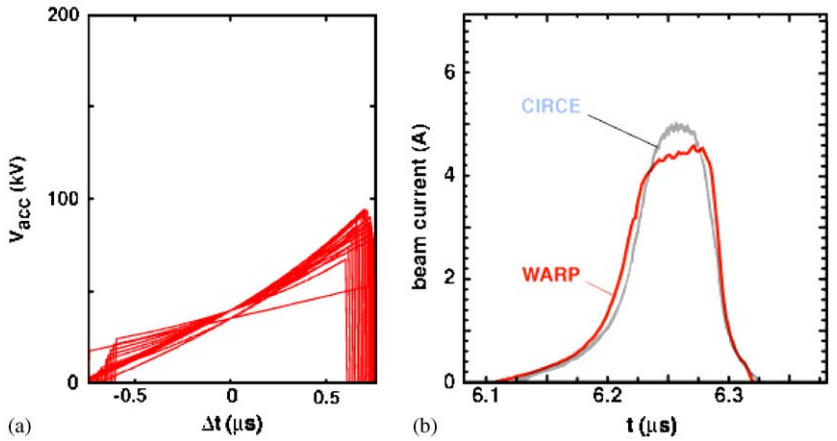

Fig. 2. Plots of (a) beam-frame shaping fields and (b) the beam current profile at the stagnation point $z = 18.4$ m when the velocity tilt is imposed by approximately triangular pulses before drift compression. The current curve labeled "CIRCE" is the same case calculated using the fluid-envelope code CIRCE.

150-MHz response frequency. These requirements substantially exceed current pulser technology. Furthermore, we expect that the high voltage and frequency content of the ear fields would make the beams extremely sensitive to any timing errors, although this sensitivity has not been tested. The other problem seen in WARP and CIRCE simulations is that the large ear fields launch a series of space-charge waves that make the final current profile highly irregular.

A third approach to achieving controlled drift compression uses a beam-shaping algorithm similar to that of Ref. [2]. Using CIRCE, a flat-topped pulse with a profile similar to the input pulse but with ten times the maximum current and a tenth the duration is first run backwards through the compression lattice, giving the velocity and current profiles that are needed at the end of the shaping section. To calculate the shaping fields, we again ignore the longitudinal space-charge field and assume instantaneous acceleration in the gaps, so that beam slices have constant velocity between gaps. Slice trajectories are then piecewise-linear curves in z–t space, with changes in slope at the induction gaps. We construct these trajectories to match the times and velocities of the input beam with the values required after shaping, and the required voltage in each gap is found from the change in the trajectory slope at that location. There are, of course, infinitely many possible trajectories that match the four boundary conditions. We have not yet developed a workable algorithm for optimizing the choice of trajectories, so here we fit a fifth-order polynomial to the initial and final data, using the additional degree of freedom to minimize the integral of the curvature along the trajectory. Piecewise-linear trajectories are then constructed by connecting the points where this polynomial intersects the gap locations. As before, ears are calculated in each gap from a CIRCE run, and the combined shaping and ear fields are tested in CIRCE by running the input beam through the entire lattice. Finally, the run is repeated using WARP3d to examine questions of phase-space dynamics, such as emittance growth.

We see from the calculated voltages in Fig. 3(a) that the choice of voltage waveforms is not optimal, since a substantial amount of energy is first removed from slices near the head and then added back later. Also, the maximum voltage exceeds what is practical in an accelerator with the short lattice period being modeled here. As expected, the current profile at stagnation obtained from CIRCE for these shaping fields, shown in Fig. 3(b), closely matches the original profile run backward through the compression lattice. The corresponding WARP3d current profile, also plotted in Fig. 3(b), shows nearly the

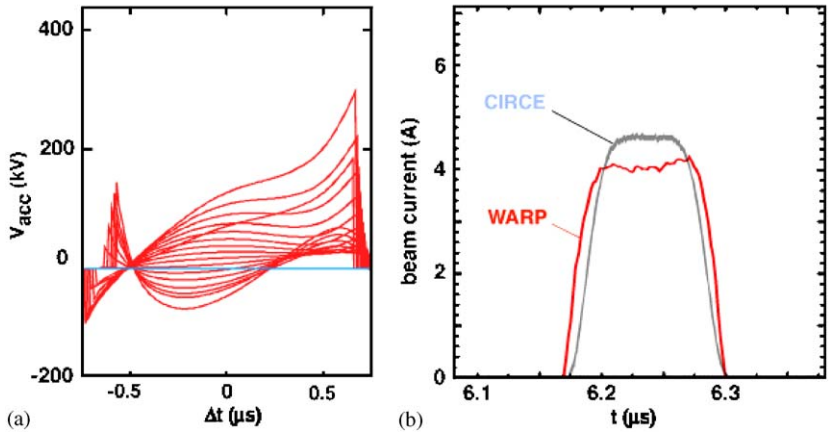

Fig. 3. Plots of (a) beam-frame shaping fields and (b) the beam current profile at the stagnation point $z = 18.4$ m when the velocity and current profiles are tailored in a shaping lattice with 17 induction cells. This particular set of shaping fields was calculated using an unoptimized algorithm that allows energy to be removed and then added back to parts of the beam. The current curve labeled "CIRCE" is the same case calculated using the fluid/envelope code CIRCE.

same 10:1 compression, and the residual variation in the average v_z along the beam is less than 2% at stagnation. However, the flat-top current is about 15% lower and broader than the intended profile. This difference between the intended and calculated current profiles may be due to using CIRCE for the backward run that gives the beam profile after shaping. As mentioned in Section 3.1, that WARP uses the correct, radially resolved space-charge field, whereas CIRCE uses the radially averaged field, so CIRCE invariably misses details of the longitudinal dynamics. We plan to test this speculation by recalculating the shaping fields using WARP3d rather than CIRCE for the backward run.

Another systematic difference between the models is that CIRCE neglects transverse emittance growth, so that beam radii calculated by CIRCE are systematically low. Due to the beam manipulation in the shaping section, WARP3d shows substantially more transverse emittance growth here than in the earlier cases with a prescribed initial tilt. However, the emittance still grows less than a factor in this case, and since the beam transverse dynamics is dominated by space charge, the resulting radius error in CIRCE is small and leads to an even smaller error in the space-charge field.

4. Conclusions

We have begun using the 3-D electrostatic PIC code WARP3d to model unneutralized drift compression of intense ion beams. WARP results for idealized runs with an initially linear head-to-tail velocity "tilt" have been with corresponding results of an analytic cold-fluid model and a more realistic fluid/envelope simulation. This comparison suggests that the analytic compression model is best used for preliminary design and scaling work. That model gives qualitatively correct compression scalings with beam perveance and initial tilt, but the predicted compression ratios are as much as 40% low for higher initial velocity tilts, due to the simple model of the longitudinal space-charge field that is used to keep the analysis tractable. The fluid/envelope CIRCE model gives more accurate estimates of compression for the modest $\Delta v/v$ values considered here, and it allows more flexibility than the analytic model in choosing the initial current, velocity, and emittance profiles. However, a particle model like WARP3d is essential both for modeling compression scenarios with larger values of perveance and initial tilt than considered here and for making credible predictions of unneutralized beam compression at HIF parameters.

Three methods have been tested for imposing a head-to-tail velocity tilt by time-dependent beam-shaping fields. Two methods prove to be unsatisfactory. Applying nearly triangular pulses achieves adequate multiplication of the peak current, but the approach allows the beam ends to expand excessively during compression and leaves a substantial velocity variation along the beam at stagnation. Using shaping fields that apply and then remove tilt before stagnation gives better control of the beam ends but places unreasonable demands on the voltage magnitude and time response of the pulsed power. The most successful of the methods constructs shaping fields that transform the input, the velocity and current profiles into those found by allowing a compressed beam to expand as it moves backward through the drift lattice. A compression ratio of 10:1 is achieved, with adequate control of the final velocity and current profiles. The transverse emittance grows by less than a factor of two during shaping and compression, and the final longitudinal velocity varies by less than 2% along the beam length at stagnation. This method, while promising, poses a challenging optimization problem. Future work will focus on developing a more effective and robust algorithm to prepare beams for drift compression.

Acknowledgments

This work was performed under the auspices of the US Department of Energy by University of California Lawrence Livermore National Laboratory and Lawrence Berkeley National Laboratory under Contracts no. W-7405-ENG-48 and DE-AC-3-76SF00098.

References

[1] D.D.-M. Ho, S.T. Brandon, E.P. Lee, Part. Accel. 35 (1991) 15.
[2] W.M. Sharp, A. Friedman, D.P. Grote, Manipulation of high-current pulses for heavy-ion fusion, AIP Conference Proceedings 391, AIP, Woodbury, New York, 1997, pp. 27–35.
[3] H. Qin, R.C. Davidson, J.J. Barnard, E.P. Lee, Drift compression and final focus options for heavy ion fusion, Proceedings of the 2004 Heavy Ion Fusion Symposium, 7–11 June 2004, Princeton, NJ, Nucl. Instr. and Meth. A, these Proceedings.
[4] W.M. Sharp, J.J. Barnard, D.P. Grote, S.M. Lund, S.S. Yu, Envelope model of beam transport in ILSE, AIP Conference Proceedings 297, AIP, Woodbury, New York, 1994, pp. 540–548.
[5] A. Friedman, D.P. Grote, I. Haber, Phys. Fluids B 4 (1992) 2203.
[6] W.M. Sharp, D.A. Callahan, D.P. Grote, Fusion Eng. Design 32 (1996) 201.
[7] W.M. Sharp, D.P. Grote, Phys. Rev. Special Topics: Accel. Beams 5 (2002) 094202.